BANK OF CANADA
BANQUE DU CANADA

Staff Working Paper/Document de travail du personnel 2017-8

Adoption Costs of Financial Innovation: Evidence from Italian ATM Cards

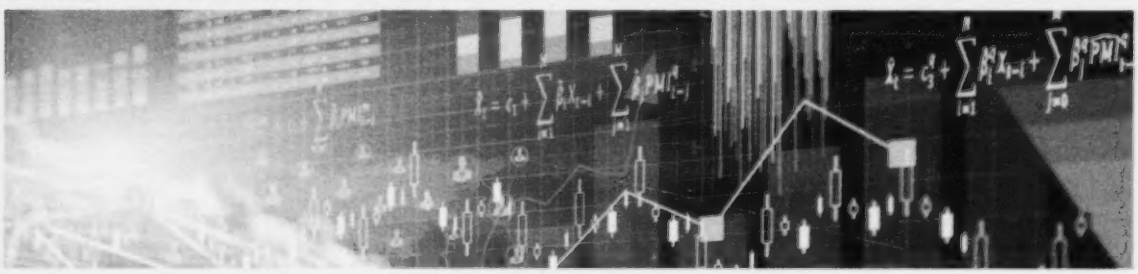

by Kim P. Huynh, Philipp Schmidt-Dengler, Gregor W. Smith and Angelika Welte

Bank of Canada staff working papers provide a forum for staff to publish work-in-progress research independently from the Bank's Governing Council. This research may support or challenge prevailing policy orthodoxy. Therefore, the views expressed in this paper are solely those of the authors and may differ from official Bank of Canada views. No responsibility for them should be attributed to the Bank.

www.bank-banque-canada.ca

Bank of Canada Staff Working Paper 2017-8

February 2017

Adoption Costs of Financial Innovation: Evidence from Italian ATM Cards

by

Kim P. Huynh,[1] Philipp Schmidt-Dengler,[2] Gregor W. Smith[3] and Angelika Welte[1]

[1]Currency Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
khuynh@bankofcanada.ca
awelte@bankofcanada.ca

[2]Department of Economics
University of Vienna
philipp.schmidt-dengler@univie.ac.at

[3]Department of Economics
Queen's University
smithgw@econ.queensu.ca

ISSN 1701-9397

© 2017 Bank of Canada

Acknowledgements

Smith acknowledges the Social Sciences and Humanities Research Council of Canada and the Bank of Canada research fellowship program for support of this research. Schmidt-Dengler acknowledges financial support from the German Science Foundation through Sonderforschungsbereich Transregio 15. Dörte Heger provided excellent research assistance. For their constructive comments, we thank Jason Allen, Thomas Lemieux, Héctor Pérez Saiz, Ben Tomlin, and seminar and conference participants. The views in this paper represent those of the authors alone and are not those of the Bank of Canada.

Abstract

The discrete choice to adopt a financial innovation affects a household's exposure to inflation and transactions costs. We model this adoption decision as being subject to an unobserved cost. Estimating the cost requires a dynamic structural model, to which we apply a conditional choice simulation estimator. A novel feature of our method is that preference parameters are estimated separately, from the Euler equations of a shopping-time model, to aid statistical efficiency. We apply this method to study ATM card adoption in the Bank of Italy's *Survey of Household Income and Wealth*. There, the implicit adoption cost is too large to be consistent with standard models of rational choice, even when sorted by age, cohort, education or region.

Bank topics: Bank notes; Econometric and statistical methods; Financial services
JEL codes: E41; D14; C35

Résumé

Le choix discret d'adopter une innovation financière influe sur l'exposition d'un ménage à l'inflation et aux coûts de transaction. Nous modélisons cette décision d'adoption en considérant un coût non observé. L'estimation du coût nécessite un modèle structurel dynamique, auquel nous appliquons un estimateur de simulation de choix conditionnel. Notre méthode est nouvelle en ce qu'elle permet d'estimer séparément les paramètres de préférences, à partir des équations d'Euler d'un modèle portant sur le temps passé à faire les magasins. Notre démarche vise à améliorer l'efficience statistique. Nous appliquons cette méthode pour étudier l'adoption d'une carte bancaire en nous appuyant sur l'enquête sur le revenu et la richesse des ménages (*Survey of Household Income and Wealth*) menée par la Banque d'Italie. D'après les données d'enquête, le coût implicite d'adoption est trop élevé pour être compatible avec les résultats des modèles standard de choix rationnel, et ce, même lorsqu'un tri est opéré selon l'âge, la cohorte, le niveau d'instruction ou la région.

Sujets : Billets de banque; Méthodes économétriques et statistiques; Services financiers
Codes JEL : E41; D14; C35

Non-technical Summary

Motivation and Question

Adopting a bank card provides ongoing benefits to consumers. For example, withdrawing banknotes at an automated teller machine (ATM) helps households use cash efficiently. However, households also may incur an implicit (unobserved), one-time adoption cost to access this new financial technology. Measuring this cost is important to predicting the speed with which such a technology can spread.

Methodology

We study ATM card adoption using the Bank of Italy's *Survey of Household Income and Wealth* (SHIW). The SHIW is a rich household panel data set spanning the period 1991–2004. It contains information about each household's characteristics, financial portfolios, and banking habits. During this period the proportion of the population who held an ATM card rose from 29 percent to 58 percent. To understand this relatively slow adoption rate for ATM cards, we estimate the unobserved adoption cost using a money-demand model in which household's use cash to economize on shopping time. The model is estimated in two stages, with the first stage describing the optimal amount of cash to be held by the consumer. In the second stage, we estimate the cost of adoption based on the first stage model and the adoption probabilities found in the data.

Key Results and Contributions

We estimate the benefits from using an ATM card to be quite large, which means that a large, implicit adoption cost is necessary to rationalize the slow adoption observed historically. We find that the overall adoption cost is about 4,000 to 18,000 euros depending on the assumptions of our statistical models. The striking finding of our research is that this adoption cost is implausibly large. Sorting it by age, education, or region does not alter this finding. It is possible that there is some difference between card-holders and non-holders that affects consumption but cannot be observed in the SHIW. But a more likely explanation seems to be consumer inertia, a challenge to the rational adoption model. One key implication is that pricing a financial innovation to encourage its adoption may be very costly. Wider initiatives to encourage upgrades in account status also may be slow to have any effect.

1 Introduction

The extensive margin or participation decision of portfolio decisions—whether to open an account or change account type—affects economists' assessments of subjects such as the inflation tax or the transactions cost borne by households. For example, when we model the impact of inflation on saving or portfolio decisions, we need to account for the impact on account holding. Investigations of stock holding, like those of Haliassos and Bertaut (1995) and Guiso, Haliassos, and Jappelli (2003), suggest that differences in participation in holding risky assets may reflect differences in a perceived cost, broadly defined to include information and fees. This cost may vary over individuals, time, or countries and in part explain changes in participation. There is a lot of heterogeneity in account-holding decisions, which should shed light on the nature of the perceived adoption cost. Moreover, real option theory suggests that a small cost may explain significant postponements in adopting new instruments for saving or financial innovations. Yet, the question of whether such a cost can plausibly explain the lack of adoption of certain accounts remains unanswered. For example, Barr (2007) reports that more than 8.4 million US households do not have either a checking or savings account.

The adoption-cost theory is not the only possible approach to account-holding decisions. As an alternative, Barberis, Huang, and Thaler (2006) provide a behavioral perspective on holding stocks. But because the nature of account holding affects so many issues in portfolio theory and asset pricing, it seems worthwhile to develop and apply methods to assess the adoption-cost perspective. The goals of this paper are to devise and apply such methods.

Estimating an adoption cost that cannot be observed directly requires a structural model. At the same time, we want to keep track of heterogeneity across households, by age, region, education, and other measures. Incorporating these realistic features makes repeated solution of a dynamic programming problem costly. Estimation here instead uses a conditional choice simulation estimator, like that introduced by Hotz, Miller, Sanders, and Smith (1994). The resulting computational tractability allows us to include many dimensions of observable heterogeneity (by age, education, location, or household composition, for example). Our application is to the decision to adopt an ATM card over the last 25 years in Italy. We draw mainly on the Bank of Italy's *Survey of Household Income and Wealth* (*SHIW*), which is the richest available survey that tracks household financial decisions.

An innovation in this paper is to combine the simulation estimator with the estimation of preference parameters via Euler equations for households with or without ATM cards. The Euler equations come from a shopping-time model that describes both the intensive margin of money holding and the additional gains from holding an ATM card. The money-demand function implied by the shopping-time model also allows for the diffusion of ATM machines and bank branches over the historical sample: It is important to control for this diffusion in

banking services in estimating the adoption cost. Drawing on this Euler equation information means that the simulation estimator estimates the adoption cost with greater precision.

Our main finding is that the implicit adoption cost, which closely reproduces actual adoption probabilities, is very large. It is always possible that there is some unobserved difference between ATM card-holders and non-card-holders that rationalizes this result. But this seems unlikely, given the detailed household information in the survey. For example, characteristics such as age, education, or regional status cannot explain it. Thus, this finding suggests consumer inertia, and poses a challenge to the rational adoption model. It implies that a small subsidy (by banks) to promote adoption may have very little effect.

Section 2 describes the planning problem we attribute to households. Section 3 briefly reviews related research on account adoption, to set our approach in context. Section 4 discusses the Bank of Italy's *SHIW* and other data sources. Section 5 outlines the econometric building blocks (including the Euler equation estimation) while section 6 describes the simulation estimator. Section 7 discusses the results from the simulation estimator for the adoption cost. Section 8 contains a summary.

2 Household Choice Problem

A household in time period $t \in \{1, \ldots, T\}$ is indexed by $i \in \{1, \ldots, N\}$. It has a vector of characteristics, denoted x_{it}, which may include the size of the household, income, and measures of education. It has one deterministically time-varying characteristic, its age (in practice, the age of the oldest labour income earner), denoted a_{it}. Its planning horizon runs to age A. Thus, it makes decisions for a horizon of $A - a_{it}$ periods, so that older households have fewer remaining decisions and shorter horizons. Households also encounter exogenous variables: the inflation rate π_t and the nominal regional deposit interest rate r_{it}.

Households base decisions on their wealth w_{it} and on their account status, labelled I_{it}, which takes a value of 1 for those with an ATM card and 0 for those without one. They decide on real consumption, c_{it}, real money holdings, m_{it}, and whether to open an account. We adopt a shopping-time model of money holding, as outlined by McCallum (1989, pp 35–41) and Walsh (2003, pp 96–100). Utility is increasing in consumption and leisure. Leisure is decreasing in shopping time which, in turn, is increasing in consumption and decreasing in real balances. The net result is that utility is increasing in consumption (with the direct, positive effect exceeding the negative effect due to shopping time) and money holdings.

The utility function is:

$$u(c_{it}, m_{it}) = (1 + \gamma I_{it})^{\omega} \frac{c_{it}^{1-\alpha} - 1}{1 - \alpha} + \frac{m_{it}^{1-\omega} - 1}{1 - \omega} d(t)^{\omega}. \tag{2.1}$$

Thus, holding an ATM card (or a similar innovation) increases the effective consumption service flow associated with a given money balance, to an extent governed by γ. The function d serves two key purposes. First, its value at a given time gives a weight to cash holding relative to consumption in the utility function. Second, this relative weight can vary over time. We thereby allow for the diffusion of bank branches and ATMs over time, developments that may increase the effectiveness of money holding. We model this function as a deterministic time trend:

$$d(t) = e^{\lambda_0 + \lambda_1 t}. \tag{2.2}$$

This curve is common to both card-holders and non-card-holders. A trend is chosen as it fits the data well and offers computational simplicity, as we do not project forward any trend into simulations beyond T. The idea is that there is a time trend in the utility effect of a given stock of money balances. This trend, d, falls over the sample. Thus, the effective consumption associated with a given money balance rises. This trend has a similar but ongoing effect whether one holds a card or not. In addition, adding an ATM card has a level effect that economizes on money or allows more consumption per unit of money holding.

Let j count across the time periods from t to $A - a_{it}$. A household plans sequences of consumption, money holding, and account status to maximize:

$$\mathrm{E} \sum_{j=0}^{A-a_{it}} \beta^k u(c_{i,t+j}, m_{i,t+j}), \tag{2.3}$$

with discount factor β, and subject to initial conditions and to transitions.

The shopping-time utility function reflects the benefit of upgrading to an ATM card. But changing account status also has a fixed cost with mean κ. There is also a random shock to the adoption cost, which is normal with mean 0 and standard deviation σ_ε. We later consider the possibility that the adoption cost may depend on household characteristics such as age, education, or location.

Households know how wealth evolves, but we do not. In most data sets we do not have detailed information on sources of investment income, fees, and transactions costs. We want our method to be applicable to data sets that include information on wealth and savings, but do not include detailed information on income sources, returns, or portfolio shares. Stango and Zinman (2009) document actual fees paid on checking accounts and credit cards for an unusual sample of 917 US households, but such tracking does not take place in any standard, large survey. Thus, we model the wealth transitions (which may depend on account status) non-parametrically.

Heterogeneity across households can enter the problem in two ways. First, it enters the transitions for wealth and consumption, which can depend on the characteristics of the

household. Second, it affects the horizon through age. This feature allows for the possibility that older households are less likely to open accounts because they have fewer periods over which the utility gains can offset the fixed cost.

A household without an account holds a perpetual American call option in that it can open an account at any time at a fixed cost. But we cannot value this option using arbitrage-free methods, because there is incomplete insurance. However, we do use an optimality property familiar from real options or from dynamic programming with a fixed cost: the difference in lifetime, expected utility values between those who hold accounts and those who do not is equal to the fixed cost of opening an account. Real option theory shows that a relatively small cost κ can lead to a significant delay. Dixit and Pindyck (1994) and Stokey (2009) demonstrate this effect in investment problems, for example. This effect occurs because the value of continuing without an account includes the option to exercise at a future time. We attempt to quantify this adoption cost.

Now that the choice problem and notation have been outlined, we next briefly summarize some related research, then turn to our application and estimation.

3 Related Research

This project is related to behavioral perspectives on financial decisions, to econometric work on the participation decision, to numerical portfolio models, and to studies of the demand for money. This section briefly sets our work within these contexts.

First, financial decisions, including the decision to adopt a new type of account, may be made infrequently and may be subject to deferred benefits that are hard to measure. Households thus may well make mistakes in these decisions, as a wealth of research in behavioral economics has emphasized. Agarwal, Driscoll, Gabaix, and Laibson (2009) define mistakes to include paying unnecessarily high fees and interest payments. They find a U-shaped pattern in financial mistakes with age. Calvet, Campbell, and Sodini (2007) study under-diversification and under-participation in risky asset markets in the Swedish population. They find that these investment errors are associated with low education and low wealth. Given this pattern, they conclude that these behaviors genuinely are mistakes and so probably are not due to heterogeneity in risk aversion or background risk. DellaVigna (2009) summarizes field evidence on departures from the standard model of economic decision-making, including non-standard preferences, non-rational expectations, and non-standard decision-making.

Behavioral economists also then address the question of whether competition may tend to correct these mistakes or whether, in addition, regulation may improve welfare. This approach might rationalize a range of regulations that restrict payday loans or that require

banks to offer certain types of low-fee accounts or open branches in specific areas (as required by the US *Community Reinvestment Act* for example). A number of institutions and economists also promote financial literacy, which usually includes encouraging saving and opening a bank account. Lusardi and Mitchell (2014) outline research on this topic.

We explore the possibility that some households may rationally not adopt an ATM card, because of the cost involved in doing so. Non-adoption may be a wise choice if fees are high or ATM locations are inconvenient. The rational-option approach to not holding an account would be perfectly consistent with an observation that a small change in information or costs—a nudge, to use Thaler and Sunstein's (2008) term—could promote account holding. We estimate what this change in costs would be, based on household characteristics. Vissing-Jørgensen (2003) suggests that showing non-participation to be rational requires that investigators find the participation cost to not be implausibly large. She concludes that relatively modest running costs can explain non-participation in US stock markets. We explore whether this conclusion holds for those without ATM cards.

Second, the extensive margin or participation choice is a key issue in empirical work on household portfolios. Most work on participation or account adoption uses limited dependent variable econometrics to statistically explain the dichotomous variable I_{it}. Miniaci and Weber (2002) provide an excellent discussion of the econometric issues in estimating these models with data from household surveys. Most empirical work concerns the decision to hold risky assets. Guiso, Haliassos, and Jappelli (2003), Perraudin and Sørensen (2000), and Vissing-Jørgensen (2002) study participation in stock markets. The same methods have also been applied to the decision to open a bank account. Attanasio, Guiso, and Jappelli (2002) study the demand for currency allowing for the adoption decision.

A number of studies use panel data that allow researchers to track adoptions and potentially allow for unobserved heterogeneity across households. Alessie, Hochguertel, and van Soest (2004), for example, track the ownership of stocks and mutual funds in a panel of Dutch households. Huynh (2007) studies the adoption of ATM cards using the *SHIW* panel. These studies typically find that previous participation is significant in statistically modelling current participation. A key issue is whether this pattern reflects true state-dependence or persistent unobserved exogenous variables. This pattern of persistence may be evidence of a fixed cost to adoption and so may help identify that cost.

Third, a number of researchers have studied the extensive margin of stock holding using numerical portfolio models solved by dynamic programming. For example, Halisassos and Michaelides (2003) introduce a fixed cost into an infinite-horizon, consumption-portfolio model. Gomes and Michaelides (2005) also solve and simulate a life-cycle model with a fixed cost of holding stocks. These studies calibrate planning problems and carefully study the outcomes.

Alan (2006) takes the important step from simulation to estimation, by indirect inference. She estimates preference parameters and a fixed cost of entry to a stock market using the solved consumption-portfolio model so that statistics from a participation equation in the simulated data match those from the same equation in historical panel data. Sanroman (2007) adopts a similar method. She outlines a planning problem that involves both a participation decision and an asset-allocation or portfolio problem. She then solves the dynamic program by discretization. Finally, she estimates parameters by indirect inference using the Italian *SHIW* with a logit model as the auxiliary estimating equation. She estimates that the participation cost for holding stocks varies from 0.175 to 6 percent of income, or from 10 to 1,126 euros, with households with higher education implicitly facing lower costs.

Our study differs from the work on dynamic consumption-portfolio models in that we do not solve a dynamic program. A fixed cost produces an inward kink in the upper envelope of the unconditional value function that makes the policy function discontinuous: a problem for numerical work. In addition, with a finite horizon, there is a set of age-dependent policy functions. And we wish to include horizons and adoption costs that may be heterogeneous across households. All these features make the repeated solution of a dynamic program costly. We adopt a conditional choice simulation estimator like those developed by Hotz, Miller, Sanders, and Smith (1994) and based on the method of Hotz and Miller (1993). It is, of course, an open question as to whether the properties of estimation would be improved in practice with estimation that involves repeated dynamic programming. Aguirregabiria and Mira (2007) provide a complete overview of estimators.

Fourth, in studying cash holding using the *SHIW*, we follow in the footsteps of Attanasio, Guiso, and Jappelli (2002), Lippi and Secchi (2009), and Alvarez and Lippi (2009). Attanasio, Guiso, and Jappelli (2002) study the demand for money using a generalized inventory model, and note the effects of ATM card usage. They then calculate the welfare cost of inflation. Lippi and Secchi (2009) show how to account for trends in the availability of banking services in order to estimate money-demand parameters. Alvarez and Lippi's (2009) study of household cash management is the state of the art in applications of the inventory-theoretic framework. They model the household's cash withdrawal conditional on the adoption of an ATM card. Their framework also measures changes over time in withdrawal costs. They find a relatively small benefit to adopting an ATM card, although they note that it is based only on a reduction in withdrawal costs and not on the card's use as a debit card. Yang and Ching (2014) model both the extensive and intensive margins, using the Baumol–Tobin model to describe the latter. They estimate a significantly larger cost of ATM card adoption. We adopt a general, shopping-time model of money holding that is also widely used in macroeconomic theory and allows for the trend and elasticity findings of Attansio, Guiso, and Jappelli (2002) and Lippi and Secchi (2009).

In sum, we combine some of the economic structure from numerical portfolio models—planning horizons and parameters of discounting and intertemporal substitution—with the ability to accommodate all the household heterogeneity and econometric tractability of the discrete-choice econometric models. We also try to ensure that our period utility function reflects recent research on the demand for money.

4 Data Sources

Our study relies mostly on household-level data from the Bank of Italy's *SHIW*. However, we require macroeconomic and aggregate data from a variety of external sources. We discuss these aggregate variables, then the *SHIW*.

4.1 Inflation and Interest Rates

We use data on inflation and interest rates from a variety of sources. The inflation rate, measured as the per annum change in consumer prices, is taken from the *International Financial Statistics* of the International Monetary Fund. The data are on an annual basis from 1989 to 2010. The Banca d'Italia *Base Informativa Pubblica* online historical database is the source for regional bank branch density and also for regional nominal deposit interest rates. These interest rates are constructed from a variety of historical tables at a quarterly frequency. The quarterly data are then aggregated to an annual frequency using simple sum averaging to derive annual data from 1989 to 2010.

4.2 Survey of Household Income and Wealth

The Italian *SHIW* is the gold standard for panel surveys involving wealth and savings. It has detailed information on account status, wealth, and consumption, and the largest and longest coverage of any such panel. The *SHIW* is the main data source for studies on money demand and financial innovation by Attanasio, Guiso, and Jappelli (2002), Alvarez and Lippi (2009), and Lippi and Secchi (2009), among others.

The *SHIW* is a biennial survey run by the Banca d'Italia. We use the 1991, 1993, 1995, 1998, and 2000, 2002, and 2004 waves. We stop at 2004 as one of our main variables—average currency holdings—is discontinued from 2006 onwards, with the exception of 2008. The three-year spacing from 1995 to 1998 was a result of the Banca d'Italia switching survey providers. The Banca d'Italia spends considerable resources to ensure that the data are nationally representative, as outlined by Brandolini and Cannari (2006). The *SHIW* survey is a rotating panel, with about 8,000 households per wave. The rotating panel design is incorporated because there is an attrition rate of roughly 50 percent. Jappelli and Pistaferri

(2000) provide an extensive discussion of the quality of the *SHIW* data and also provide a comparison with Italian National Accounts data to address issues of sample representativeness, attrition, and measurement. Details about the variables we used are available in a separate technical appendix.

ATM cards involve a small annual fee, but no additional charges for withdrawals at machines owned by the issuing bank. Their first benefit is that they allow card-holders to withdraw cash rapidly and when banks are closed. Checking accounts bear interest, so the ability to make withdrawals at lower cost can reduce foregone interest earnings from holding cash. A second benefit is that they can be used as point-of-sale debit cards for retail transactions. Despite these benefits, though, the use of cash remained very widespread in Italy throughout this period.

Table 1 reveals that the fraction of households with an ATM card in 1991 was 29 percent and that it steadily increased to 58 percent in 2004. Given the attrition rate in the survey, one might wonder how many actual ATM card adoptions are observed: there are many. On average, the share of households that did not have an ATM card in the previous wave of the survey, were in both the current and previous waves, and had a card in a given, current wave was 16.7 percent.

Next, we focus on average currency holdings, consumption, and wealth. All the nominal variables are expressed in 2004 equivalent euros. During this period, the average currency holdings fell for households both with and without an ATM card. However, with the exception of 1991, the average cash holdings of ATM card-holders were lower than those of ATM non-card-holders. Not surprisingly, those with ATM cards tended to have higher consumption and financial wealth than those without ATM cards. Notice that the difference in consumption and wealth increased over time, as detailed by Jappelli and Pistaferri (2000).

5 Econometric Building Blocks

Our simulation estimator takes as inputs several statistical building blocks. First, we describe transitions for exogenous variables such as the regional interest rate and the inflation rate. These transition functions are denoted g. Second, we estimate the parameters of the money-demand function implied by the maximization in section 2. Third, we estimate transitions for endogenous variables—consumption, wealth, and ATM card adoption—denoting these transitions f. The next three subsections describe these steps in turn.

5.1 Interest Rate and Inflation Process

The inflation rate π_t is the year-to-year growth rate of the consumer price index (CPI), from 1989 to 2010. Figure 1 shows the distribution of the deposit interest rate (denoted r_t) across participants in the survey, for each wave. Both the level and dispersion of nominal deposit interests were highest in the early part of the 1990s. The average nominal deposit interest rate ranged from 8.6 percent in 1991 to a low of 0.40 percent in 2004 with a peak of 8.9 percent in 1993.

To parametrize the transition function for $\{\pi_t, r_t\}$, we use ordered vector autoregressions (VARs) and test the lag length with standard information criteria. We penalize models with large numbers of parameters given the short time-series sample. Let t count years (not two-year periods). We have annual data but we need two-year transition functions to align with the *SHIW*, so it makes sense to estimate those directly. We work with natural logarithms to guarantee positive, simulated interest rates and inflation rates.

We find that inflation can be described autonomously:

$$\ln \pi_t = a_0 + a_1 \ln \pi_{t-2} + \epsilon_{\pi t}, \tag{5.1}$$

with $\epsilon_{\pi t} \sim IID(0, \sigma_\pi^2)$. In each region the deposit rate is well described by:

$$\ln r_t = b_0 + b_1 \ln r_{t-2} + b_2 \ln \pi_t + \epsilon_{rt}, \tag{5.2}$$

with $\epsilon_{rt} \sim IID(0, \sigma_r^2)$.

This setup ensures that $cov(\epsilon_{\pi t}, \epsilon_{rt}) = 0$ (which simplifies simulations). We use this specific ordering because it fits with the difference in the time periods to which the inflation rate and interest rate in a given year apply. The inflation rate in year t measures the growth rate in consumer prices from year $t-1$ to t, while the deposit rate measures the interest rate that applies from t to $t+1$.

We estimate the r-equation for each of 20 regions and report the average estimates over this set (rather than averaging the interest rates, which would lead to an understatement of uncertainty in a typical region). In practice, though, the variation in estimates across regions is quite small, so we do not record a subscript for the household or region in the interest rate process.

Table 2 contains the estimates for the parameters, their standard errors, and the two residual variances. Our simulator later makes $\{\epsilon_{\pi t}, \epsilon_{rt}\}$ jointly normal and, with this ordering, the two shocks are uncorrelated. We shall denote the transition functions for inflation and the deposit interest rate by g_π and g_r, respectively.

5.2 Intratemporal Euler Equations

We have not found empirical work that estimates parameters of the shopping-time model, despite its appearance in textbooks and theoretical work, perhaps because of the shortage of data sets that track all of household consumption, leisure, real wages, and money holdings. Our new feature is to allow for holding an ATM card, or a similar innovation, to reduce shopping time and hence increase the consumption associated with a given bank balance. It turns out that incorporating that feature allows identification of both the relative role of money in the utility function and of the effect of ATM card-holding.

Consumption expenditures, c, measure real non-durable consumption, and money holdings, m, measure average currency. I is an indicator with $I = 1$ for ATM card-holders and $I = 0$ otherwise. Over the last 25 years in Italy, the nominal interest rate has trended down while consumption has trended up. According to standard models of money demand—with a positive consumption elasticity and negative interest rate elasticity—money holding should have trended up for both reasons. Yet the *SHIW* data in Table 1 show that money holding has trended *down*. This was true for both adopters and non-adopters.

To scale the decline in money holding, we construct the ratio mr/c (the interest rate–weighted money-consumption ratio) and graph it in Figure 3 for 33,591 household-year observations for non-card-holders and then for 21,936 observations for card-holders. Notice that in each category, there is a downward trend over time. The same pattern is evident with additional controls such as the age of the head of household. Figure 3 uses the same vertical scale for both categories, so it also illustrates that card-holders tend to have lower money balances at given values of consumption and interest rates.

Simultaneously, adoption has increased over time. So, we need to be careful not to attribute the fall in the ratio of mr/c entirely to ATM card adoption for, if we do, we will find an unrealistically large benefit to adoption and hence also overestimate the cost in order to fit the fact that adoption was incomplete. As section 2 noted, we interpret this pattern for both groups as evidence of the diffusion of improvements in the supply of banking services. Hester, Calcagnini, and de Bonis (2001) document changes in the ATM location practices of Italian banks between 1991 and 1995. They suggest that banking deregulation in Italy circa 1990 led to an increase in national banking in Italy. The increase in competition in turn led to an increase in the number of branches and ATMs after deregulation.

Alvarez and Lippi (2009, Table 2) show that there was an increase in the density of bank branches and ATMs during 1993–2004. We confirm this fact by plotting the diffusion of regional bank branch density for the period 1991–2004. Figure 2 displays this density in the five regions of Italy: Northeast (NE), Northwest (NW), Centre (C), South (S), and Islands. The overall trend is positive. There are regional variations, though: the Northeast had the highest banking concentration throughout the sample, closely followed by the Northwest and

Centre regions. The South and the Islands were less developed banking areas.

Both ATM card-holders and non-card-holders should have benefited from this banking deregulation. Alvarez and Lippi (2009, Table 1) shows that the number of withdrawals for both ATM card-holders and ATM non-card-holders increased during this period. Also, the level of currency fell for both ATM card-holders and ATM non-card-holders users, while m/c was relatively constant, so that mr/c should fall as the interest rate fell, just as we showed in Figure 1. Lippi and Secchi (2009) show conclusively that controlling for the diffusion of banking services—measured, in their case, by bank branches per capita—is necessary to obtain accurate money-demand elasticities.

To capture these patterns, the period utility function for household i at time t combines equations (2.1) and (2.2):

$$u(c_{it}, m_{it}) = (1 + \gamma_i I_{it})^{\omega} \frac{c_{it}^{1-\alpha} - 1}{1-\alpha} + \frac{m_{it}^{1-\omega} - 1}{1-\omega} \left(e^{\lambda_{0,i} + \lambda_1 t}\right)^{\omega}, \tag{5.3}$$

where the exponential function describes the diffusion of banking improvements. This functional form captures the key feature of the shopping-time specification. To account for heterogeneity, we let γ_i and $\lambda_{0,i}$ vary by household i. When $\gamma_i > 0$, holding an ATM card increases the effective bank balance, or equivalently increases the effective flow of consumption spending that can be financed from a given account.

The time trend acts as proxy for the diffusion of banking services over time. An alternative might entail using the number of bank branches or other indicators, although those would vary by household location. Note that this diffusion applies to both ATM card-holders and non-holders: We do not simply assume a diffusion of ATM cards but model that endogenously, based on such factors as consumption and the interest rate opportunity cost.

Estimation is based on the static or intratemporal, first-order condition. The derivatives of the utility function with respect to consumption and money are:

$$u_c = (1 + \gamma_i)^{\omega} c_{it}^{-\alpha}, \tag{5.4}$$

$$u_m = m_{it}^{-\omega} \left(e^{\lambda_{0,i} + \lambda_1 t}\right)^{\omega}. \tag{5.5}$$

The Euler equation is then $u_m = r_t u_c$, or in money-demand form:

$$m_{it} = \frac{e^{\lambda_{0,i} + \lambda_1 t}}{1 + \gamma_i I_{it}} \cdot \frac{c_{it}^{\alpha/\omega}}{r_t^{1/\omega}}. \tag{5.6}$$

Thus money holding is increasing in the transactions variable c_{it}, decreasing in the holding cost variable r_t, and decreasing in ATM card status. Holding an ATM card allows households to economize on real balances to an extent measured by γ_i. Equivalently, the level difference between the two panels of Figure 3 identifies γ_i. Notice also that the ratio mr/c can trend down, due to the diffusion function, as it does in the survey data.

To estimate the parameters of this money-demand equation, we first take natural logarithms of (5.6) and attach an error term η_{it} as follows:

$$\ln(m_{it}) = \lambda_{0,i} + \lambda_1 t - \ln(1+\gamma_i) I_{it} + \frac{\alpha}{\omega}\ln(c_{it}) - \frac{1}{\omega}\ln(r_{it}) + \eta_{it}. \tag{5.7}$$

Further, we assume that $\lambda_{0,i}$ and γ_i are functions of the education level, region, and birthyear of the household head:

$$\lambda_{0,i} = \lambda_{0,R_i} + \lambda_{0,S_i} + \lambda_{0,Y}\text{birthyear}_i + \lambda_{0,Y2}\text{birthyear}_i^2, \tag{5.8}$$

$$\gamma_{0,i} = \gamma_{0,R_i} + \gamma_{0,S_i} + \gamma_{0,Y}\text{birthyear}_i + \gamma_{0,Y2}\text{birthyear}_i^2, \tag{5.9}$$

where R_1 is the Northwest and Northeast regions, R_2 the Centre region, R_3 the South and Islands regions, S_1 is education level of at most primary; S_2 is education level of at least some secondary and no post-secondary, and S_3 is education level with some post-secondary education. The birthyear can be calculated from the age of the respondent and the survey year. The birthyear obtained in this manner is then normalized and scaled by subtracting the earliest birthyear among the survey respondents (1875) and dividing the result by 100. The resulting normalized birthyear variable takes values between 0 and 1.11. We also normalize the time variable t by subtracting 1989, the earliest wave of the survey data. Because of the linear relationships $\sum R_i = \sum S_i = 1$, a baseline category must be chosen for the regression. We choose the baseline category to be a college-educated resident of central Italy in the year 1989. The parameters $\lambda_{0,i}$, λ_1, γ_i, α/ω and $1/\omega$ in (5.7) are then estimated by linear regression with robust standard errors that are clustered within households.

Figure 4 represents the heterogeneity of the $\lambda_{0,i}$ and $\gamma_{0,i}$ in terms of region, birthyear, and education. Most of the variation in the $\lambda_{0,i}$ or the intercept of the Euler equation is due to the region, as the value for the Northwest and Northeast is much lower than those for the Centre, South, and Islands. This finding reflects lower money holdings in the North relative to the other regions. For $\gamma_{0,i}$, most of the variation is due to education and birth cohort. Those with higher education (tertiary) have a lower $\gamma_{0,i}$ in all the regions. There is a monotonic increase in the $\gamma_{0,i}$ with the birth cohort.

The last column of Table 3 shows the parameter estimates for (5.7). We find that the parameters of the utility function vary with place of residence, education, birthyear, and ATM card ownership. The consumption elasticity is 0.38 and the interest elasticity is 0.17; these arise from utility exponents of $\hat{\alpha} = 2.64$ and $\hat{\omega} = 5.77$. The resulting average value of $\hat{\gamma}_i$ is 0.26. A time trend is estimated via the diffusion coefficient $\hat{\lambda}_1$; that is, the coefficient for $year - 1989$, which is negative and statistically significant.

The special case of unit elasticities arises when $\alpha = \omega = 1$. In this case, the utility function is expressed in natural logarithms on consumption and cash holdings as

$$\lim_{\alpha\to 1}\lim_{\omega\to 1} u(c,m) = (1+\gamma_i)\ln(c_{it}) + \ln(m_{it})\left(e^{\lambda_{0,i}+\lambda_1 t}\right), \tag{5.10}$$

and the Euler equation in logarithms simplifies to:

$$\ln(m_{it}) = \lambda_{0,i} + \lambda_1 t - \ln(1+\gamma_i) I_{it} + \ln(c_{it}) - \ln(r_{it}) + \eta_{it} \tag{5.11}$$

or

$$\ln\left(\frac{m_{it} r_{it}}{c_{it}}\right) = \lambda_{0,i} + \lambda_1 t - \ln(1+\gamma_i) I_{it} + \eta_{it}. \tag{5.12}$$

The central column of Table 3 shows the parameter estimates for this special case (5.12). In this case, the resulting average value of $\hat{\gamma}_i$ is 0.46. And again there is a significant, negative time trend. Figure 5 shows the heterogeneity in the estimates $\lambda_{0,i}$ and $\gamma_{0,i}$ in terms of region, birthyear, and education. Most of the variation in $\lambda_{0,i}$, the intercept of the Euler equation, is again due to the region as the value in the North is much lower than in the Centre, South, and Islands. For $\gamma_{0,i}$, most of the variation is due to birth cohort. The estimate of $\gamma_{0,i}$ has a quadratic relationship with the birth cohort with a slight decrease then a large increase. There is not much variation due to education, as was found in the Constant Relative Risk Aversion (CRRA) case in Figure 4.

Overall, we find that the estimates from the more general specification in the last column show that these unit-elasticity restrictions are rejected. But we display the results for two reasons. First, our description of money demand was motivated in part by the patterns in the ratio mr/c shown in Figure 3. Table 3 confirms that the downward trend in that figure and the difference between card-holders and non-holders are statistically significant. Second, we shall later estimate adoption costs with both functional forms, to illustrate the effect of the money-demand function on the cost estimates.

5.3 Consumption, Wealth, and Adoption Processes

We next estimate the reduced-form transitions of consumption and wealth conditional on whether a household adopts an ATM card ($I = 1$) or not ($I = 0$). Time is annual and is indexed by t, but the survey is sampled every two years (three years between 1995 and 1998). Therefore, the lag of a variable represents a two-year spacing.

Transitions may depend on the age of household i, age a_{it}, time-invariant household characteristics, x_i, nominal deposit interest rates, r_{it}, the inflation rate, π_t, income, y_{it}, and wealth w_{it}. For simplicity, we collect these state variables in a vector z_{it}:

$$z_{it} \equiv \{w_{it}, a_{it}, x_{it}, y_{it}, r_{it}, \pi_t\}. \tag{5.13}$$

The transitions of consumption (c_{it}) and wealth (w_{it}) are then of the form:

$$c_{it} = f_c(z_{it-2}|I_{it}). \tag{5.14}$$

$$w_{it} = f_w(z_{it-2}|I_{it}) \tag{5.15}$$

There is substantial heterogeneity in consumption and wealth between ATM card-holders and ATM non-card-holders. We clean the data by dropping households with negative consumption (nine households). Next, we drop the 1st and 99th percentiles of wealth and consumption. The estimates are computed using ordinary least squares that are sample-weighted and use robust standard errors. Attempts at using fractional polynomials did not yield a large improvement in fit relative to the added complexity. The results are available upon request. All the monetary and financial variables are deflated to 2004 euro-equivalent numbers.

Table 4 contains the results of consumption transition function (5.14). The covariates included are: wealth, real regional deposit interest rates, age profile, gender, employment indicators, education indicators, household size measures (adults and children), residence indicators, and a set of year and region dummy variables. The overall fit of the consumption profile (pooled) is about 53.4 percent with the ATM non-card-holder case having a higher fit of 51.8 percent relative to the 35.6 percent for ATM card-holders. Most of the variables are statistically and economically significant. There is statistical difference across the groups in the coefficients of net disposable income and wealth.

The wealth transition results are displayed in Table 5. Note that consumption and income are both excluded from the regression. The overall fit of the model measured by R^2 is 57.9 percent, with 56.9 percent, for ATM card-holders and 55.7 percent, for ATM non-card-holders. The lag of wealth is significant and statistically different between ATM and non-ATM holders.

Finally, we also need transitions for ATM card-holder status itself. We estimate transitions for all households pooled, for those who previously held an ATM card ($I_{i,t-2} = 1$), and for those who previously did not hold an ATM card ($I_{i,t-2} = 0$). These ATM card transition results are displayed in Table 6. The key component is in the third column: ATM card adoption policy, estimated using a probit model. That case is labelled as follows:

$$Pr[I_{it} = 1 | I_{i,t-2} = 0, z_{it}] = f_I(z_{it}). \tag{5.16}$$

Notice that account adoption is statistically related to a variety of household characteristics, including age and education, as well as to the interest rate. A high interest rate is associated with greater ATM card adoption. Households with less education or located in rural regions are less likely to adopt. This dependence will serve to identify the adoption cost. (We also studied this problem using a logit instead of probit specification. The results do not quantitatively change. These robustness checks are available upon request.)

5.4 Tracking Money and Wealth Changes at Adoption

Our estimates of the costs of adoption will depend on the estimated benefits of adoption, which in turn depend on the estimated shift in utility measured by γ. In subsection 5.2 we controlled for a range of fixed effects that may differ across households, to ensure that we did not overestimate γ by pooling dissimilar households. However, the reader might wonder whether there is evidence that adoption is directly associated with changes in money holding.

Figure 6 plots the money-consumption (mr/c) and wealth-consumption (w/c) ratios over sequences of three waves of the *SHIW* for the adopters, denoted by (0,1,1), the always-adopters, denoted by (1,1,1), and the never-adopters, denoted by (0,0,0). The plots apply to three time windows: 1991–1993–1995 (denoted W1), 1998–2000–2002 (denoted W2), and 2000–2002–2004 (denoted W3).

The top panel shows the ratio mr/c. It illustrates that the never-adopters have the highest ratios, followed by adopters, and then the always-adopters. In each window, the ratio mr/c is decreasing, consistent with earlier analysis that the overall mr/c ratio is falling over time. The bottom panel shows the w/c ratio for the same households.

Comparing the three groups in the top panel suggests that adoption per se is associated with a fall in money holding relative to consumption (in time periods W1 and W2), an economizing on money balances, which will raise utility. Comparing the same three groups in the bottom panel suggests that adoption is associated with a rise in financial wealth relative to consumption (in W1 and W3 but not W2).

This descriptive analysis of these ratios during ATM card adoption is instructive but we note that we are dealing with non-experimental or observational data. Therefore, we cannot treat adoption of ATM cards as a randomly assigned treatment for which we can compare outcomes. Instead, we need to model the decision to adopt an ATM card via our structural model. In the next section, we use the econometric building blocks from this section to construct a simulation-based estimator.

6 Adoption Cost Estimation

Using the econometric building blocks, we can now construct a simulation-based estimator for the adoption cost. The roots of such an estimator were discussed by Hotz, Miller, Sanders, and Smith (1994) and Pakes (1994). The focus is on estimating the mean cost of adoption κ and the standard deviation σ_z of the adoption cost shock. ATM card adoption is modelled as a finite-horizon, discrete, dynamic choice with a terminal state (ATM card adoption).

We begin with a key feature of the *SHIW*: We observe data only in biennial intervals.

Thus, observations are in years that are elements of the set

$$T_O = \{1989, 1991, 1993, 1995, 1998, 2000, 2002, 2004\};$$

i.e. if $t \in T_O$, we refer to a period in which data are observed. We thus assume that each household makes a consumption decision c_{it} and a cash-holding decision m_{it} and receives utility from consumption only every other period. Similarly, the household decides on ATM card adoption I_{it} only in periods $t \in T_O$ and every two years after that. The adoption decision is irreversible: Once an ATM card has been adopted, the household keeps it. Thus, the adoption indicator I_{it} is weakly increasing in time t: $I_{it} \in (I_{i,t-2}, 1)$. We study households that have a bank account and age above 25. A household's age is given by a_{it}, which can reach at most $A(= 100)$. We also assume that retirement occurs for everyone at age 65. The index j counts years within simulations.

The estimation algorithm begins by simulating paths for consumption, money holding, and wealth, under a scenario in which the household adopts in period t and also under the scenarios in which the household does not adopt in period t, but adopts in periods $t + j$ for $j \in \{2, 4 \ldots, 2 \cdot \lfloor (A - a_{it})/2 \rfloor\}$ or never adopts during the lifetime of the household head. The aim is to obtain simulated "choice-specific values" for the choice of non-adoption and adoption in year j. Tildes denote simulated variables. Simulated variables for the path generated with no adoption in period t have superscript 0 while the path with adoption has superscript 1. Along the 0-path, a household may choose to adopt an ATM card at a later stage, but once it has adopted, it cannot un-adopt the card. Since the adoption state is terminal, there are only finitely many possible paths for a household that does not adopt in period t, namely $\lfloor (A - a_{it})/2 \rfloor + 1$ corresponding to adoption in periods $t + j$ for $j \in \{2, 4 \ldots, 2 \cdot \lfloor (A - a_{it})/2 \rfloor\}$ and the path where the household does not adopt during its lifetime. If the household adopts in time period t, only one path exists.

The conditional value function v_{it}^I is defined as the present discounted value of choosing I in period t (net of the adoption shock) and then behaving optimally in future periods $t + j$. The conditional value function under adoption in period t is:

$$\begin{aligned} \tilde{v}_{it}^1 &= u(\tilde{c}_{it}^1, \tilde{m}_{it}^1; I_{it} = 1) - \kappa + \beta^2 \sum_{j=2,4}^{2 \cdot \lfloor (A - a_{it})/2 \rfloor} \beta^{j-2} \cdot E\left(u(\tilde{c}_{it+j}^1, \tilde{m}_{it+j}^1; I_{it+j} = 1, \tilde{z}_{i,t-2}^1\right) \\ &= u(\tilde{c}_{it}^1, \tilde{m}_{it}^1; I_{it} = 1) - \kappa + \beta^2 \cdot (\tilde{V}_{i,t+2}^1). \end{aligned}$$

In this expression, $u(\tilde{c}_{it}^1, \tilde{m}_{it}^1; I_{it} = 1) - \kappa$ is the current period utility payoff in period t net of the adoption cost and $\tilde{V}_{i,t+2}^1 = \beta^2 \sum_{j=2,4}^{2 \cdot \lfloor (A - a_{it})/2 \rfloor} \beta^{j-2} \cdot \left(u(\tilde{c}_{it+j}^1, \tilde{m}_{it+j}^1; I_{it+j-2} = 1, \tilde{z}_{it+j-2})\right)$ is the discounted future value of lifetime utility following adoption in period t. The error terms of the utility contributions after adoption are equal to 0, since the household is not facing

any decision problem whence:

$$\tilde{v}_{it}^1 = u(\tilde{c}_{it}^1, \tilde{m}_{it}^1; I_{it} = 1) - \kappa + \beta^2 \sum_{j=2,4}^{2\cdot\lfloor (A-a_{it})/2 \rfloor} \beta^{j-2} \cdot \left(u(\tilde{c}_{it+j}^1, \tilde{m}_{it+j}^1; I_{it+j} = 1, \tilde{z}_{i,t+j-2}^1)\right).$$

If no adoption takes place in period t, then:

$$\tilde{v}_t^0 = u(\tilde{c}_{it}^0, \tilde{m}_{it}^0, I_{it} = 0) + \beta^2 \cdot (\tilde{V}_{i,t+2}^0 | I_{it} = 0, z_{i,t}^0).$$

We now explain how to calculate the expected lifetime utility $\tilde{V}_{i,t+2}^0$. The key observation is the identity:

$$\begin{aligned}
&(\tilde{V}_{i,t+j}^0 | I_{t+j-2} = 0, \tilde{z}_{it+j-2}^0) \\
&= (1 - f_I(\tilde{z}_{it+j}^0) \cdot \left(u(\tilde{c}_{it+j}^0, \tilde{m}_{it+j}^0; I_{it+j} = 0, \tilde{z}_{t+j-2}^0) + \beta^2 \cdot (\tilde{V}_{i,t+j+2}^0 | I_{i,t+j} = 0, \tilde{z}_{it+j}^0)\right) \qquad (6.1) \\
&\quad + f_I(\tilde{z}_{it+j}^0) \cdot \Big(u(\tilde{c}_{it+j}^0, \tilde{m}_{it+j}^0; I_{it+j} = 1, z_{t+j-2}^0) - \kappa - \sigma_\epsilon \cdot E(\epsilon | \epsilon < \Phi^{-1}(f_I(\tilde{z}_{it+j}^0))) \\
&\qquad + \beta^2 \cdot (\tilde{V}_{i,t+j+2}^0 | I_{t+j} = 1, \tilde{z}_{it+j}^0)\Big),
\end{aligned}$$

where $f_I(\tilde{z}_{it+j})$ denotes the probability that the household adopts in time period $t+j$, conditional on not having adopted yet. If the adoption cost shocks are standard normal, the conditional expectation is given by:

$$E(\epsilon | \epsilon < \Phi^{-1}(f_I(\tilde{z}_{it+j})) = -\frac{\varphi(\Phi^{-1}(f_I(\tilde{z}_{it+j}))}{f_I(\tilde{z}_{it+j})},$$

where φ and Φ are the standard normal probability and cumulative distribution functions, respectively. Note that the state variable $\tilde{z}_{i,t+j}^I$ in period $t+j$ depends on $I_{t+j'}$ and $\tilde{z}_{it+j'}^I$ for $j' < j$. Once adoption has taken place, the value function evolves deterministically.

As a result, the expected future value along the path with adoption after j time periods is defined as:

$$\begin{aligned}
\tilde{V}_{i,t+2}^0(j) &:= (\tilde{V}_{i,t+2}^0; I_{i,t} = 0, \ldots, I_{i,t+j-2} = 0, I_{i,t+j} = 1) \\
&= \sum_{k=0,2,4}^{j-2} \beta^k (u(\tilde{c}_{it+k}^0, \tilde{m}_{it+k}^0, I_{it+k} = 0) + \\
\beta^j &\left((u(\tilde{c}_{it+j}^0, \tilde{m}_{it+j}^0, I_{it+j} = 1)) - \kappa - \sigma_\epsilon \cdot E(\epsilon | \Phi^{-1}(\epsilon < f_I(\tilde{z}_{it+j}^0)))\right) + \\
&\sum_{k=j+2,j+4}^{2\cdot\lfloor (A-a_{it})/2 \rfloor} \beta^k (u(\tilde{c}_{it+k}^0, \tilde{m}_{it+k}^0, I_{it+k} = 1)).
\end{aligned}$$

If the household does not adopt in its lifetime, the expected future value is given by:

$$\tilde{V}_{i,t+2,\infty}^0 = \sum_{k=0,2,4}^{2\cdot\lfloor (A-a_{it})/2 \rfloor} \beta^k (u(\tilde{c}_{it+k}^1, \tilde{m}_{it+k}^0; I_{it+k} = 0)).$$

Using the period-wise probabilities, the probability that the household does not adopt in periods $t, \ldots, t+j-2$, but adopts in period $t+j$ can be calculated:

$$\tilde{p}_{t+j} = \begin{cases} \left(\prod_{k=2}^{j-2}(1 - f_I(\tilde{z}_{it+k}))\right) f_I(\tilde{z}_{it+j}), & a_t + j \leq A \\ 0 & a_t + j > A. \end{cases}$$

Note that the state vector $\tilde{z}_{it+k}$ depends on the household's history. The probability that the household does not adopt during its lifetime is (by abuse of notation):

$$\tilde{p}_{t+\infty} = 1 - \sum_{j=2,4}^{2\cdot[(A-a_{it})/2]} \tilde{p}_{t+j}.$$

We can now calculate:

$$\tilde{v}^0_{it} = u(\tilde{c}^0_{it}, \tilde{m}^0_{it}; I_{it} = 0) + \sum_{j=2,4}^{2\cdot[(A-a_{it})/2]} \left(\tilde{p}_{t+j} \cdot \tilde{V}^0_{i,t+2}(j)\right) + (\tilde{p}_{t+\infty} \cdot \tilde{V}^0_{i,t+2,\infty}).$$

To produce these simulated paths we use the building blocks from section 5. First, the inflation rate and interest rate are simulated using draws from the processes in Table 2. Thus, the inflation rate is from:

$$\tilde{\pi}_{t+j} = g_\pi(\tilde{\pi}_{t+j-2}) = \exp\left[0.182 + 0.694 \ln \tilde{\pi}_{t+j-2} + \tilde{\epsilon}_{\pi t}\right], \tag{6.2}$$

and the interest rate is from:

$$\tilde{r}_{t+j} = g_r(\tilde{r}_{t+j-2}, \tilde{\pi}_{t+j}) = \exp\left[-1.185 + 0.826 \ln \tilde{r}_{t+j-2} + 0.947 \ln \tilde{\pi}_{t+j} + \tilde{\epsilon}_{rt}\right] \tag{6.3}$$

with $\sigma^2_\pi = 0.162$ and $\sigma^2_r = 0.314$.

Second, for a simulated value of consumption, along with interest rates and adoption status, money holdings can be simulated from the Euler equations (5.7) or with log utility (5.12). Third, consumption itself can be simulated from the estimated transition function f_c (5.15) and wealth from the transition function f_w (5.16) given adoption state $I_{i,t+j}$. Notice that the transition functions are for the logarithms of the variables. This was motivated by the observation that the value function is linear in the logarithms of wealth, consumption, and the interest rate.

Just as in the historical data, in simulations we label the state variables in a vector $z_{i,t+j}$:

$$z_{i,t+j} \equiv \{w_{i,t+j}, a_{i,t+j}, x_{i,t+j}, r_{i,t+j}, \pi_{t+j}\}. \tag{6.4}$$

The conditional choice simulation estimator works by matching simulated and empirical choice probabilities. The latter are found from the estimated ATM card adoption policy, given in Table 6:

$$Pr[I_{it} = 1 | I_{i,t-2} = 0, z_{it}] = f_I(z_{it}). \tag{6.5}$$

The key outcomes from this algorithm are unbiased, simulated values of the choice-specific values $\tilde{v}_{it}^0$ and $\tilde{v}_{it}^1$. We obtain S simulated values $\left(\tilde{v}_{it,s}^0, \tilde{v}_{it,s}^1\right)$. We obtain our simulator by averaging over simulations:

$$v_{it}^{0,S} = \frac{1}{S}\sum_{s=1}^{S} \tilde{v}_{it,s}^0 \quad \text{and} \quad v_{it}^{1,S} = \frac{1}{S}\sum_{s=1}^{S} \tilde{v}_{it,s}^1$$

We estimate our model applying the insight from Hotz, Miller, Sanders, and Smith (1994). Observe that a household will be indifferent between adopting and not adopting whenever

$$v_{it}^1 - \sigma_\epsilon \epsilon_{it} = v_{it}^0.$$

The household adopts for all $\epsilon_{it} < \left(v_{it}^1 - v_{it}^0\right)/\sigma_\epsilon$. It thus follows that for all (i,t):

$$\Phi^{-1}\left(Pr(I_{it} = 1 | I_{i,t-2} = 0, z_{it})\right) = \left(v_{it}^1 - v_{it}^0\right)/\sigma_\epsilon.$$

Our estimator will be based on this equality, replacing the probability $Pr(I_{it} = 1 | I_{i,t-2} = 0, z_{it})$ on the left-hand side with our estimate $f_I(z_{it})$, and the choice-specific values with their simulated counterparts $(v_{it}^{0,S}, v_{it}^{1,S})$. Specifically, we estimate κ and σ_ϵ by non-linear least squares.

7 Adoption Cost Estimates and Interpretation

Table 7 summarizes the simulation and estimation for the adoption cost $\hat{\kappa}$. We estimate the adoption cost both for log utility and for CRRA utility. The discount factor is 0.95 for both utility functions. The overall estimates of $\hat{\kappa}$ (with standard errors in brackets) are 32.5 (0.17) and 21.3 (0.34) for the log utility and CRRA cases, respectively. However, the reader is cautioned to not compare these two quantities directly. For example, the estimated standard deviations of the shock to the adoption probability, $\hat{\sigma}_\epsilon$ (again with standard errors in brackets), are 0.42 (0.024) and 0.15 (0.016) for log and CRRA utility, respectively. Therefore, it is not surprising that the $\hat{\kappa}$ with CRRA utility is smaller than with log utility. Another important consideration is that these $\hat{\kappa}$ must be scaled relative to consumption, age, education, and regional profiles.

To account for this observed heterogeneity, we compute the adoption costs for 27 cells for the three variables: birth cohort (COH), defined as the oldest, middle, and young, denoted as 1, 2, and 3; education (EDU), defined as primary or less (1), secondary (2), post-secondary education (3); and regions (REG), defined as: North (1), Centre (2), South and Islands (3). We find that for each case, the parameter estimate of fixed cost ($\hat{\kappa}$) is heterogeneous among the various demographic profiles. These differences are statistically significant given the

small standard errors (not shown). The coefficient $\acute{\kappa}$ is increasing with successive cohorts and education levels while the Centre has the highest $\acute{\kappa}$ relative to the South and then North. This result seems counterintuitive since we would expect the fixed cost parameter to be lowest for the young and educated in the North, where there is highest amount of financial development. However, one should view the $\acute{\kappa}$ as the parameter that leaves the household indifferent between adopting and not adopting conditional on their utility. Therefore, we need to translate the utility into a metric that we can compare.

The estimated parameter $\acute{\kappa}$ is not directly interpretable. Therefore, to understand the adoption costs, we use a measure of how much consumption is required to make a household indifferent between adopting and not adopting. This interpretation of the compensated cost is similar to the measure employed by Cooley and Hansen (1989) to understand the welfare costs of inflation. Recall that households that have an ATM card are denoted with $I = 1$ while those that do not are denoted with $I = 0$. Denote by y the consumer's income, which does not change after the policy is introduced. The indirect utility is denoted by $V^I(y)$; that is, the maximum utility under policy I, given income y. The compensating consumption Δc is defined implicitly as $V^1(y) = V^0(y - \Delta c)$.

Subscripts i for the individual consumers are suppressed. For simplicity we will assume $y = c$. Note that $V^1 - V^0 = \sigma_\epsilon \Phi^{-1}(f_1(z))$ as in section 6, where $f_1(z)$ is the probability of adopting an ATM card at the beginning of the time period. So,

$$\sigma_\epsilon \Phi^{-1}(f_1(z)) = V^0(c - \Delta c) - V^0(c).$$

We assume that the adoption cost is paid at time period T_0, so that the difference on the right-hand-side (RHS) is $u_0^0(c - \Delta c) - u^0(c)$ and we must solve:

$$\sigma_\epsilon \Phi^{-1}(f_1(z)) = u_0^0(c - \Delta c) - u_0^0(c).$$

Note that κ does not appear because the cost is partially offset by the benefit from adoption due to $\gamma > 0$. The value Δc takes these benefits into account. If a consumer is indifferent between adopting and non-adopting after factoring in the cost and benefits of adoption, then $\sigma_\epsilon \Phi^{-1}(f_1(z)) = 0$ and $\Delta c = 0$. If $f_1(z) > 0.5$, then $\Phi^{-1}(f_1(z))) > 0$, $\Delta c < 0$ and the consumer has negative adoption cost. Conversely, if $\Phi^{-1}(f_1(z))) < 0$, the consumer has a positive adoption cost.

In solving for Δc, we ignore the utility from cash holding because it is small compared with the utility from consumption. For brevity, let $\bar{\epsilon} = \Phi^{-1}(f_1(z))$ so that:

1. **CRRA utility**: Let $u_0(c) = \frac{c^{\alpha-1}-1}{\alpha-1}$, $\alpha > 1$. In this case , $\Delta c = c - (c^{1-\alpha} + \bar{\epsilon}(1 - \alpha))^{1/(\alpha-1)}$.

2. **Log utility**: Let $u_0(c) = \ln(c)$ then $\Delta c = c(1 - e^{\epsilon})$.

From the estimates in section 5, the coefficient of relative risk aversion is $\alpha = 5.77$. Therefore, for the CRRA utility function, Δc is well defined if and only if $c^{1-\alpha} + \bar{c}(1-\alpha) > 0$. Δc is set to 0 whenever this inequality is violated. Note that $c^{1-\alpha} + \bar{c}(1-\alpha) \leq 0$ implies that $\bar{c} > 0$ hence, that $f_1(z) > 0.5$ and that the household is more likely to adopt than not to adopt. Such households should have $\Delta c < 0$. Setting $\Delta c = 0$ for them therefore overestimates the average adoption cost for their cell in Table 7.

Table 7 gives the values for Δc for the 27 cells sorted by birth cohort, education, and region. The overall Δc values are €4,429 and €17,870 for the log and CRRA utility cases, respectively. The Δc values are all larger for CRRA than log utility. There is clearly variation across age cohorts, education levels, and regions, but the values are quite large. Sorting households according to their regional trends in banking diffusion (as in section 5) does not affect this conclusion.

The large adoption cost (Δc) explains the slow pace of adoption but it also implies that the estimated benefits of adoption must also be large. Figure 3 showed that cash holding was significantly lower for ATM card-holders. Over the historical sample, one of the benefits of card holding, in higher interest income on bank account balances, fell because the interest rate itself fell, as shown in Figure 1. But the diffusion of ATMs and bank branches, captured by the time path of d, had an offsetting effect, increasing the benefits of holding an ATM card.

Our conclusion is that explaining the slow adoption and non-adoption of ATM cards requires an unrealistically large fixed cost. As Vissing-Jorgenson (2003) noted, this seems like a rejection of the rational adoption model. Methodologically, this finding shows the benefits of the two-step method that incorporates information from the Euler equation. That method leads to precision in estimating $\hat{\kappa}$. And it disciplines the search for parameter values by requiring that the utility-function parameters fit the evidence in subsection 5.2. We then use these parameters to compute the adoption costs. The finding of an implausible cost means that either (a) the utility function is mis-specified (so that ATM card-holding is correlated with some other, unobserved feature that explains the consumption and money-holding differences) or (b) there is some other explanation for slow adoption, such as a lack of information or household inertia.

In reaching these conclusions, we calibrated the discount factor as $\beta = 0.95$, but the findings are not sensitive to this value. Figure 7 graphs $\hat{\sigma}_\epsilon$ (in green, on the left vertical axis) and $\hat{\kappa}$ (in purple, on the right vertical axis). Notice that $\hat{\kappa}$ rises with β, as one would expect: The greater the weight placed on the future, the greater the benefits of adoption, and so the larger the value of the fixed cost necessary to explain non-adoption. But notice that $\hat{\kappa}$ is large for any reasonable value of β.

8 Conclusion

We have studied a dynamic, discrete choice problem in which households may adopt a banking innovation: in this case, an ATM card in Italy. We used a conditional choice simulation estimator, like that introduced by Hotz, Miller, Sanders, and Smith (1994). A key feature of the economic environment is the return or utility function. That is based on a shopping-time model of money demand, with two distinctive features. First, it allows for a gradual diffusion of bank branches and ATM machines between 1989 and 2004, which enhanced the efficiency of money holding (and so reduced the ratio of money to consumption) for both card-holders and non-card-holders. Second, it includes a parameter (γ) that isolates the additional degree to which card-holders economized on money holding.

We estimate these features of money demand via the Euler equations in a first step, using data from more than 52,000 household-year observations. We also estimate transitions for consumption and wealth for both groups. These econometric building blocks then allow quite precise estimation of the adoption cost using the simulation estimator. Our method is applicable to a range of additional financial adoption decisions.

Our striking finding is that the adoption cost is implausibly large. Sorting it by age, education, or region does not alter this finding. It is possible that there is some difference between card-holders and non-holders that affects consumption but cannot be observed in the *SHIW*. But a more likely explanation seems to be consumer inertia, a challenge to the rational adoption model. One key implication is that pricing a financial innovation to encourage its adoption may be very costly. Wider initiatives to encourage upgrades in account status may also be slow to have any effect.

References

Agarwal, Sumit, John C. Driscoll, Xavier Gabaix, and David Laibson (2009) The age of reason: Financial decisions over the lifecycle and implications for regulation. *Brookings Papers on Economic Activity* 40(2) (Fall)), 51–117.

Aguirregabiria, Victor and Pedro Mira (2007) Sequential estimation of dynamic discrete games. *Econometrica* 75, 1–53.

Alan, Sule (2006) Entry costs and stock market participation over the life cycle. *Review of Economic Dynamics* 9, 588–611.

Alessie, Rob, Stefan Hochguertel, and Arthur van Soest (2004) Ownership of stocks and mutual funds: A panel data analysis. *Review of Economics and Statistics* 86, 783–796.

Alvarez, Fernando and Francesco Lippi (2009) Financial innovation and the transactions

demand for cash. *Econometrica* 77, 363–402.

Attanasio, Orazio R., Luigi Guiso, and Tullio Jappelli (2002) The demand for money, financial innovation, and the welfare cost of inflation: An analysis with household data. *Journal of Political Economy* 110, 317–351.

Barberis, Nicholas, Ming Huang, and Richard Thaler (2006) Individual preferences, monetary gambles, and stock market participation: A case for narrow framing. *American Economic Review* 96, 1069–1090.

Barr, Michael S. (2007) An inclusive, progressive national savings and financial services policy, *Harvard Law and Policy Review* 1, 161–184.

Brandolini, Andrea and Luigi Cannari (2006) Methodological Appendix: The Bank of Italy's Survey of Household Income and Wealth. In A. Ando, L. Guiso, and I. Visco, eds. *Saving and the Accumulation of Wealth: Essays on Italian Household and Government Saving Behavior.* Cambridge University Press.

Calvet, Laurent E., John Y. Campbell, and Paolo Sodini (2007) Down or out: Assessing the welfare costs of household investment mistakes. *Journal of Political Economy* 115, 707–747.

Cooley, Thomas F. and Gary D. Hansen (1989) The inflation tax in a real business cycle model. *American Economic Review*, 79:4, 733–48.

DellaVigna, Stefano (2009) Psychology and economics: Evidence from the field. *Journal of Economic Literature* 47:2, 315–372.

Dixit, Avinash K. and Robert S. Pindyck (1994) *Investment under Uncertainty.* Princeton University Press.

Gomes, Francisco and Alexander Michaelides (2005) Optimal life-cycle asset allocation: Understanding the empirical evidence. *Journal of Finance* 60, 869–904.

Guiso, Luigi, Michael Haliassos, and Tullio Jappelli (2003) Household stockholding in Europe: Where do we stand and where do we go? *Economic Policy* 18(36) (2003) 123–170.

Haliassos, Michael and Carol C. Bertaut (1995) Why do so few hold stocks? *Economic Journal* 105, 1110–1129.

Haliassos, Michael and Alexander Michaelides (2003) Portfolio choice and liquidity constraints. *International Economic Review* 44, 143–177.

Hester, Donald, D., Giorgio Calcagnini, and Riccardo de Bonis (2001) Competition through innovation: ATMs in Italian banks. *Rivista Italiana degli Economist* (3), 359–381.

Hotz, V. Joseph and Robert A. Miller (1993) Conditional choice probabilities and the estimation of dynamic models. *Review of Economic Studies* 60, 497–531.

Hotz, V. Joseph, Robert A. Miller, Seth Sanders, and Jeffrey Smith (1994) A simulation estimator for dynamic models of discrete choice. *Review of Economic Studies* 61, 265–289.

Huynh, Kim P. (2007) Financial innovation and the persistence of the extensive margin. Mimeo: http://pages.iu.edu/~kphuynh/rsch/extmarg.html.

Jappelli, Tullio and Luigi Pistaferri (2000) The dynamics of household wealth accumulation in Italy. *Fiscal Studies* 21, 269–295.

Lippi, Francesco and Alessandro Secchi (2009) Technological change and the households' demand for currency. *Journal of Monetary Economics* 56, 222–230.

Lusardi, Annamaria, and Olivia S. Mitchell (2014) The economic importance of financial literacy: Theory and evidence. *Journal of Economic Literature* 52(1): 5–44.

McCallum, Bennett T. (1989) *Monetary Economics.* Macmillan.

Miniaci, Raffaele and Guglielmo Weber (2002) Econometric issues in the estimation of household portfolios. Chapter 4 in L. Guiso, M. Haliassos, and T. Jappelli, eds. *Household Portfolios.* MIT Press.

Pakes, Ariel (1994) Dynamic structural models, problems and prospects: mixed continuous discrete controls and market interactions, in: C. A. Sims (Ed.), *Advances in Econometrics*, Vol. 2, Cambridge University Press, 171–274.

Perraudin, William and Bent E. Sørensen (2000) The demand for risky assets: Sample selection and household portfolios. *Journal of Econometrics* 97, 117–144.

Sanroman, Graciela (2007) Estimating heterogeneous costs of participation in the risky asset markets. mimeo, CEMFI.

Stango, Victor and Jonathan Zinman (2009) What do consumers really pay on their checking and credit card accounts? Explicit, implicit, and avoidable costs. *American Economic Review* 99, 424–429.

Stokey, Nancy L. (2009) *The Economics of Inaction.* Princeton University Press.

Thaler, Richard and Cass Sunstein (2008) *Nudge: Improving Decisions about Health, Wealth, and Happiness.* Yale University Press.

Vissing-Jørgensen, Annette (2002) Towards an explanation of household portfolio choice heterogeneity: Non-financial income and participation cost structures, mimeo, University of Chicago.

Vissing-Jørgensen, Annette (2003) Perspectives on behavioral finance: Does 'irrationality' disappear with wealth? Evidence from expectations and actions. *NBER Macroeconomics Annual* 18, 139–194.

Walsh, Carl E. (2003) *Monetary Theory and Policy*. 2nd Edition, MIT Press.

Yang, Botao and Andrew T. Ching (2014) Dynamics of consumer adoption of financial innovation: The case of ATM cards. *Management Science* 60, 903–922.

Table 1: Descriptive Statistics

	1991	1993	1995	1998	2000	2002	2004
Fraction with an ATM card %	29%	34%	40%	49%	52%	56%	58%
Average currency holdings (m)							
with ATM card	741	527	398	421	374	349	341
without ATM card	696	607	457	498	438	443	458
Non-durable consumption (c)							
with ATM card	28,459	26,520	26,431	27,250	24,464	25,765	24,941
without ATM card	20,550	17,759	17,243	17,304	15,343	15,392	15,138
Financial wealth							
with ATM card	161,113	189,770	210,882	203,397	190,513	194,880	201,501
without ATM card	115,437	119,208	126,172	120,246	111,717	112,562	110,951
Interest rate or r	8.6%	8.9%	7.0%	3.2%	2.2%	1.6%	0.4%
mr/c ratios							
with ATM card	0.23	0.19	0.15	0.12	0.06	0.03	0.03
without ATM card	0.32	0.33	0.26	0.22	0.10	0.07	0.05
Observations	7916	7770	7840	6794	7616	7660	7639

Note: Average currency holdings, consumption, and wealth are expressed in terms of €2004. The source is the Bank of Italy's *Survey of Household Income and Wealth*.

Table 2: Interest Rate and Inflation Process

$$
\begin{aligned}
\ln \pi_t &= a_0 + a_1 \pi_{t-2} + \epsilon_{\pi t} \\
\ln r_t &= b_0 + b_1 \ln r_{t-2} + b_2 \ln \pi_t + \epsilon_{rt} \\
& \epsilon_{\pi t} \sim IIN(0, \sigma_\pi^2) \\
& \epsilon_{rt} \sim IIN(0, \sigma_r^2)
\end{aligned}
$$

Parameter	Estimate	Standard Error
a_0	0.182	0.256
a_1	0.694	0.208
σ_π^2	0.162	
b_0	-1.185	0.288
b_1	0.826	0.127
b_2	0.947	0.312
σ_r^2	0.314	

Notes: Estimation uses annual observations from 1989–2010 on CPI inflation and deposit rates averaged across regions. Deposit rates are from the *Banca d'Italia Base Informativa Pubblica* online historical database.

Table 3: Intratemporal Euler Equation

	$\ln(mr/c)$	$\ln(m)$
$\ln(r)$		-0.1733***
		[0.0075]
$\ln(c)$		0.3786***
		[0.0084]
$I_{it}\times$ primary	0.0451**	0.1364***
	[0.0224]	[0.0193]
$I_{it}\times$ secondary	0.0469**	0.0842***
	[0.0204]	[0.0175]
$I_{it}\times$ North	0.0351*	0.0770***
	[0.0203]	[0.0177]
$I_{it}\times$ South and Islands	0.1270***	0.1602***
	[0.0228]	[0.0199]
$I_{it}\times$ birthyear	1.5153***	-0.3829
	[0.4554]	[0.3913]
$I_{it}\times$ birthyear squared	-1.4174***	0.0135
	[0.3291]	[0.2818]
ATM card	-0.7138***	-0.1243
	[0.1570]	[0.1356]
North	-0.2795***	-0.2402***
	[0.0142]	[0.0123]
South and Islands	0.0903***	0.0616***
	[0.0141]	[0.0121]
Primary	0.2436***	-0.0288**
	[0.0151]	[0.0130]
Secondary	0.1595***	0.0079
	[0.0157]	[0.0133]
Normalized birthyear	-2.5251***	0.9310***
	[0.2604]	[0.2221]
Normalized birthyear squared	1.7200***	-0.7851***
	[0.2061]	[0.1751]
Year: 1989	-0.2182***	-0.0660***
	[0.0008]	[0.0016]
Constant	-4.7579***	2.0164***
	[0.0812]	[0.0972]
Adjusted R^2	0.6826	0.1830
LogL	-64620	-56499
Observations	52081	52081

Note: Standard errors are in brackets and *, **, *** denote statistical significance at the 10, 5, and 1 percent level. All variables are deflated to €2004. The base demographic profile is college educated, born in 1875, living in the Centre of Italy in the year 1989, not having an ATM card.

Table 4: Consumption Transition

	All	ATM Card	No ATM Card
Wealth	0.0892***	0.0969***	0.0785***
	[0.0015]	[0.0027]	[0.0017]
Real regional interest rates	-0.6503***	0.0684	-0.0241
	[0.1159]	[0.1843]	[0.1512]
Age	0.0086***	0.0064***	0.0071***
	[0.0011]	[0.0020]	[0.0014]
Age squared	-0.0001***	-0.0001***	-0.0001***
	[0.0000]	[0.0000]	[0.0000]
Male	0.0526***	0.0244***	0.0715***
	[0.0059]	[0.0091]	[0.0074]
Employed	0.1009***	0.0704***	0.1194***
	[0.0073]	[0.0111]	[0.0093]
Self-employed	0.1049***	0.0652***	0.1518***
	[0.0087]	[0.0129]	[0.0112]
No schooling	-0.4775***	-0.3462***	-0.4642***
	[0.0134]	[0.0363]	[0.0197]
Elementary school	-0.4151***	-0.3386***	-0.4070***
	[0.0105]	[0.0144]	[0.0179]
Middle school	-0.3056***	-0.2545***	-0.3198***
	[0.0104]	[0.0127]	[0.0182]
High school	-0.1704***	-0.1486***	-0.1914***
	[0.0102]	[0.0117]	[0.0187]
Northwest	0.0141**	0.0084	-0.0096
	[0.0071]	[0.0101]	[0.0095]
Northeast	0.0212***	0.0144	0.0049
	[0.0072]	[0.0101]	[0.0097]
South	-0.2440***	-0.2006***	-0.2292***
	[0.0072]	[0.0125]	[0.0084]
Islands	-0.2532***	-0.2005***	-0.2471***
	[0.0092]	[0.0162]	[0.0104]
Constant	8.5457***	8.5757***	8.6084***
	[0.0345]	[0.0546]	[0.0458]
RMSE	0.3498	0.3495	0.3410
R^2	0.5343	0.3563	0.5181
Adjusted-R^2	0.5341	0.3558	0.5179
LogL	-21908	-9183	-11824
Observations	59476	24994	34482

Note: Standard errors are in brackets and *, **, *** denote statistical significance at the 10, 5, and 1 percent level. Year and region dummy variables are suppressed for brevity. All variables are deflated to €2004. The base demographic profile is male, college educated, not in the labour force and living in an urban area of the Centre of Italy in the year 1989.

Table 5: Wealth Transition

	All	ATM Card	No ATM Card
Wealth previous period	0.6714***	0.6463***	0.6773***
	[0.0110]	[0.0158]	[0.0145]
Real regional interest rates	-0.7920	0.4260	-0.4649
	[0.5001]	[0.6655]	[0.7972]
Age	0.0283***	0.0203**	0.0276***
	[0.0062]	[0.0092]	[0.0086]
Age squared	-0.0002***	-0.0001	-0.0002**
	[0.0001]	[0.0001]	[0.0001]
Male	0.0917***	0.0216	0.1358***
	[0.0272]	[0.0380]	[0.0382]
Employed	-0.0760**	-0.1162***	-0.0384
	[0.0324]	[0.0401]	[0.0511]
Self-employed	0.2267***	0.1807***	0.3018***
	[0.0348]	[0.0505]	[0.0479]
No schooling	-0.6714***	-0.6547***	-0.5860***
	[0.0615]	[0.2416]	[0.0882]
Elementary school	-0.4979***	-0.4172***	-0.4561***
	[0.0412]	[0.0563]	[0.0753]
Middle school	-0.3437***	-0.3317***	-0.3160***
	[0.0409]	[0.0509]	[0.0789]
High school	-0.1549***	-0.1420***	-0.1686**
	[0.0378]	[0.0439]	[0.0807]
Northwest	-0.0574	-0.0867**	-0.0742
	[0.0358]	[0.0440]	[0.0555]
Northeast	0.0348	0.0235	0.0054
	[0.0351]	[0.0440]	[0.0532]
South	-0.1703***	-0.0960*	-0.1561***
	[0.0366]	[0.0549]	[0.0495]
Islands	-0.0500	-0.0419	-0.0166
	[0.0410]	[0.0579]	[0.0555]
Constant	3.0037***	3.5412***	2.8153***
	[0.1915]	[0.2897]	[0.2724]
RMSE	1.0611	0.9347	1.1486
R^2	0.5786	0.5695	0.5567
Adjusted-R^2	0.5781	0.5685	0.5558
LogL	-29039	-12297	-16413
Observations	19652	9107	10545

Note: Standard errors are in brackets and *, **, *** denote statistical significance at the 10, 5, and 1 percent level. Year and region dummy variables are suppressed for brevity. All variables are deflated to €2004. The base demographic profile is male, college educated, not in the labour force and living in an urban area of the Centre of Italy in the year 1989.

Table 6: ATM Transition

	All	ATM Card	No ATM Card
Age	0.0485***	0.0217	0.0357***
	[0.0049]	[0.0156]	[0.0119]
Age squared	-0.0006***	-0.0003**	-0.0004***
	[0.0000]	[0.0001]	[0.0001]
Male	0.0392	0.1530**	0.0678
	[0.0256]	[0.0682]	[0.0538]
Employed	0.0810***	0.0676	0.0768
	[0.0288]	[0.0879]	[0.0653]
Self-employed	-0.1356***	-0.3966***	-0.0843
	[0.0346]	[0.1020]	[0.0780]
No schooling	-1.6449***	-1.4927***	-1.1139***
	[0.0625]	[0.1729]	[0.1349]
Elementary school	-1.1982***	-0.8591***	-0.8409***
	[0.0414]	[0.1240]	[0.0968]
Middle school	-0.7061***	-0.6344***	-0.5060***
	[0.0399]	[0.1162]	[0.0954]
High school	-0.2271***	-0.4011***	-0.0808
	[0.0391]	[0.1139]	[0.0971]
Northwest	0.3850***	0.2934***	0.3204***
	[0.0303]	[0.0804]	[0.0708]
Northeast	0.3247***	0.4251***	0.2254***
	[0.0318]	[0.0816]	[0.0694]
South	-0.7342***	-0.4025***	-0.4144***
	[0.0317]	[0.0866]	[0.0611]
Islands	-0.5563***	0.0032	-0.4579***
	[0.0389]	[0.1203]	[0.0821]
Real regional interest rates	-17.2415***	-3.6498***	-7.6724***
	[0.5079]	[1.3665]	[1.0566]
Constant	-0.4256***	0.9987**	-1.1252***
	[0.1289]	[0.4337]	[0.3346]
LogL	-30386	-3005	-4786
Observations	61228	8106	10820

Note: Standard errors are in brackets and *, **, *** denote statistical significance at the 10, 5, and 1 percent level. Year and region dummy variables are suppressed for brevity. All variables are deflated to €2004. The base demographic profile is male, college educated, not in the labour force and living in an urban area of the Centre of Italy in the year 1989.

Table 7: Adoption Cost Structural Estimates

					Log Utility		CRRA Utility	
CELL	COH	EDU	REG	$\bar{c}^*$	$\hat{\kappa}$	Δc	$\hat{\kappa}$	Δc
1	1	1	1	15,224	29.6	4,906	6.7	15,220
2	1	1	3	16,694	34.1	5,748	15.2	16,690
3	1	1	4	13,741	18.7	6,736	1.2	13,738
4	1	3	1	18,568	29.1	4,671	12.8	18,561
5	1	3	3	19,460	33.6	5,990	24.9	13,619
6	1	3	4	17,466	18.1	7,045	5.1	17,462
7	1	4	1	24,872	35.0	1,821	26.0	5,707
8	1	4	3	25,642	40.4	3,503	45.1	7,680
9	1	4	4	25,411	23.3	6,463	13.0	25,404
10	2	1	1	21,151	31.7	4,772	13.4	21,145
11	2	1	3	22,249	36.7	6,567	25.8	22,244
12	2	1	4	17,321	20.6	6,709	5.3	17,318
13	2	3	1	22,601	32.1	3,142	21.9	22,592
14	2	3	3	25,194	37.3	5,060	39.4	25,188
15	2	3	4	20,464	20.9	6,172	10.6	20,459
16	2	4	1	28,643	38.4	(690)	41.4	24,957
17	2	4	3	30,631	44.4	1,057	70.9	30,614
18	2	4	4	25,917	26.3	3,938	22.6	25,908
19	3	1	1	20,002	42.2	3,575	20.6	19,996
20	3	1	3	24,597	47.4	5,555	36.6	24,592
21	3	1	4	17,480	31.1	5,501	10.0	17,476
22	3	3	1	20,458	47.8	2,126	34.7	20,449
23	3	3	3	23,032	52.7	4,204	58.3	23,025
24	3	3	4	15,707	34.1	4,308	18.0	15,702
25	3	4	1	23,196	55.2	(1,190)	62.0	—
26	3	4	3	27,227	61.3	(476)	101.6	27,211
27	3	4	4	21,118	39.7	3,320	34.1	21,110
Overall				19,584	32.5	4,429	21.3	17,870

Note: The discount factor is 0.95 for both utility functions. The estimated standard deviation of the shock to the adoption probability, $\hat{\sigma}_\epsilon$, is 0.42 and 0.15 for log and CRRA utility, respectively. The $\hat{\kappa}$ and adoption costs (Δc) are computed for 27 cells for three age cohorts (COH): 1 (birth year of 1936 or earlier), 2 (birth year of 1937 to 1950), and 3 (birth year of 1951 or later); three education categories (EDU): 1 (primary or less), 3 (secondary), and 4 (post-secondary and above); and three regional groups (Reg): 1 (North), 3 (Centre), and 4 (South and Islands). The within-cell mean non-durable consumption ($\bar{c}^*$) and Δc are are provided in €2004. Brackets contain quantities estimated to be negative. An em-dash (—) means that choice probabilities could not be inverted.

Figure 1: Nominal Deposit Interest Rates

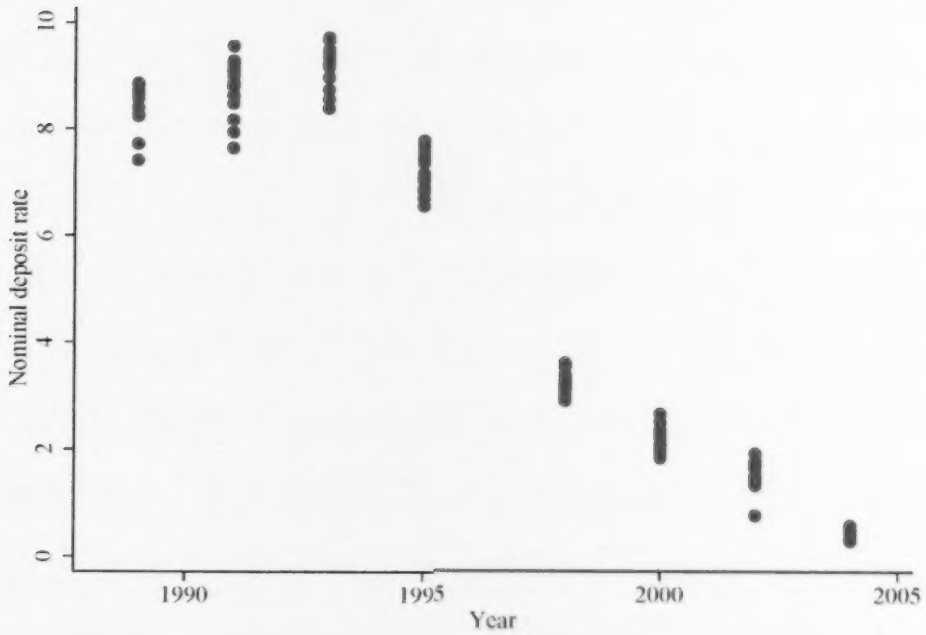

Note: Regional nominal bank deposit interest rates for households in the Banca d'Italia *Base Informativa Pubblica* online historical database.

Figure 2: Regional Bank Branch Density

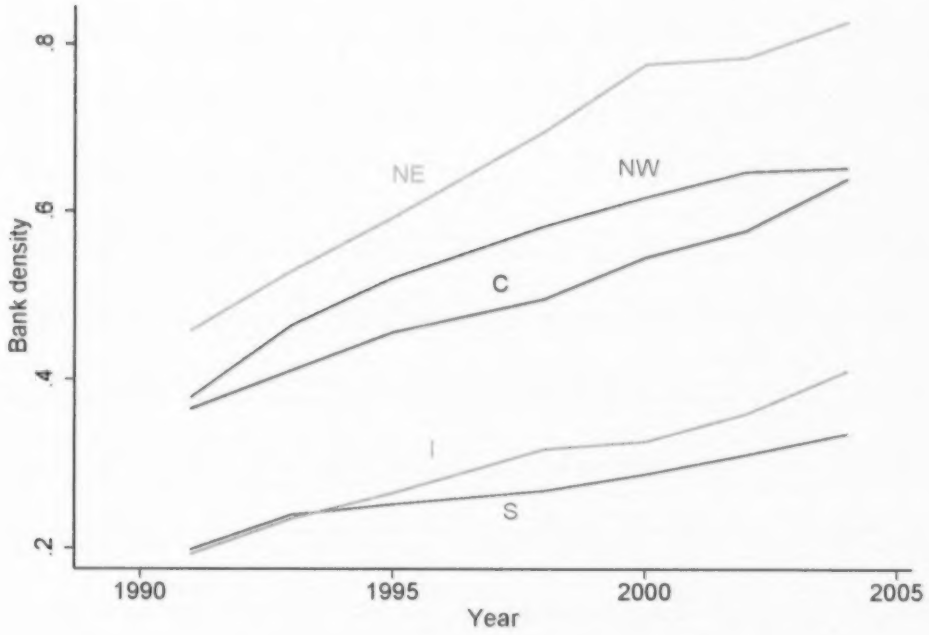

Note: Authors' calculations of the regional bank branch density are taken from Lippi and Secchi (2009), variable name *bank_pop_city*. This graph displays the regional bank branch density in the five regions of Italy: Northeast (NE), Northwest (NW), Centre (C), South (S), and Islands (I).

Figure 3: Money-Consumption Ratio for ATM Card-Holders and ATM Non-Card-Holders

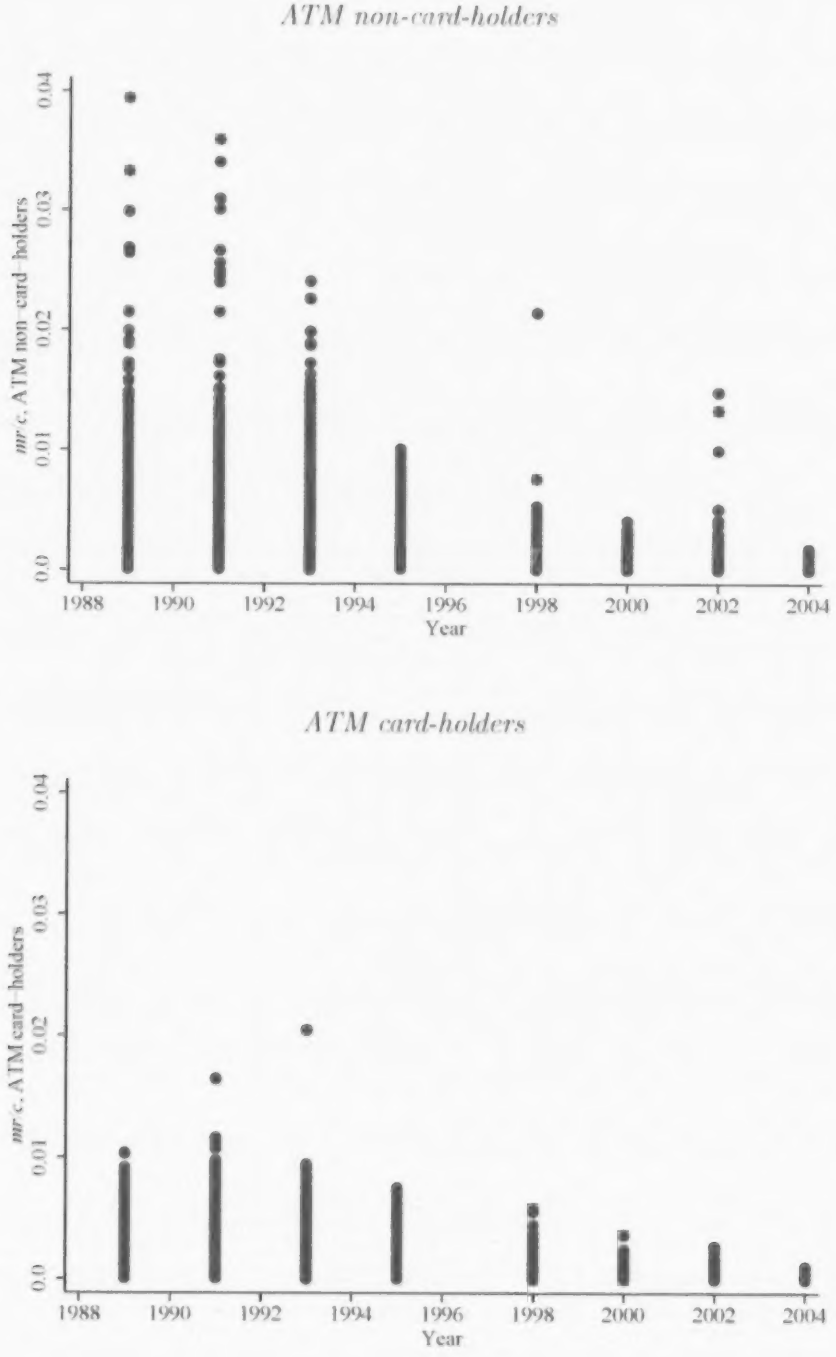

Note: Authors' calculations of the money-consumption ratio (mr/c) for ATM card-holders and ATM non-card-holders from the *SHIW*.

Figure 4: CRRA Euler Equation

$$\lambda_{0,i} = \lambda_{0,R_i} + \lambda_{0,S_i} + \lambda_{0,Y}\text{birthyear}_i + \lambda_{0,Y2}\text{birthyear}_i^2$$
$$\gamma_{0,i} = \gamma_{0,R_i} + \gamma_{0,S_i} + \gamma_{0,Y}\text{birthyear}_i + \gamma_{0,Y2}\text{birthyear}_i^2$$

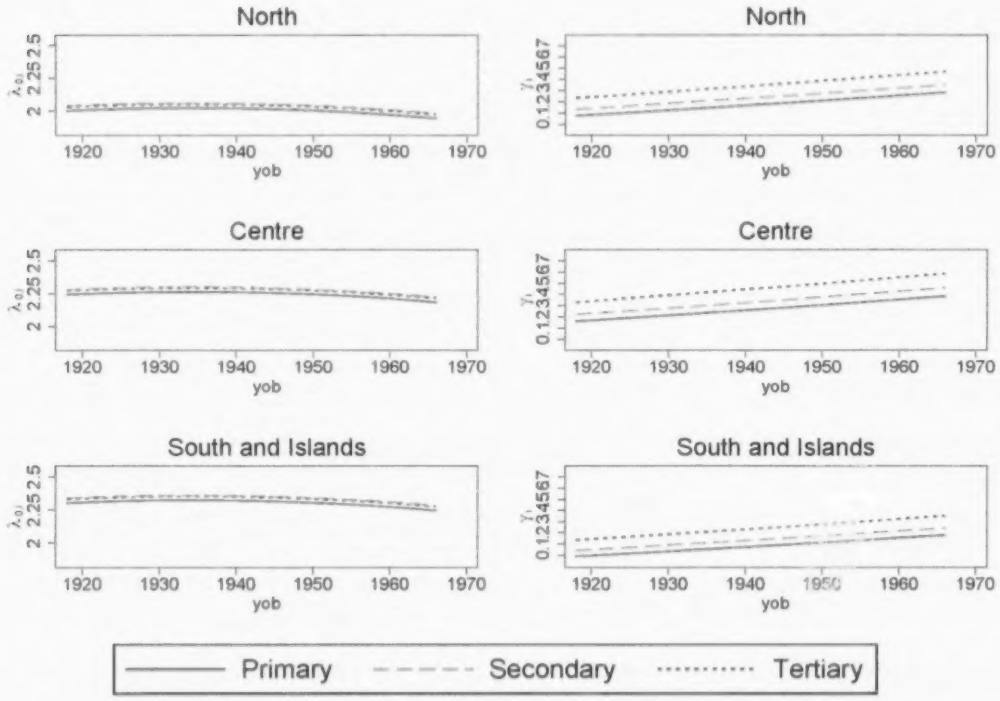

Notes: These figures are generated from Table 3. They show the estimated heterogeneity in the level ($\hat{\lambda}_{0,i}$) and ATM card-holders shift ($\hat{\gamma}_{0,i}$) in money demand.

Figure 5: Log Euler Equation

$$\lambda_{0,i} = \lambda_{0,R_i} + \lambda_{0,S_i} + \lambda_{0,Y}\text{birthyear}_i + \lambda_{0,Y2}\text{birthyear}_i^2$$
$$\gamma_{0,i} = \gamma_{0,R_i} + \gamma_{0,S_i} + \gamma_{0,Y}\text{birthyear}_i + \gamma_{0,Y2}\text{birthyear}_i^2$$

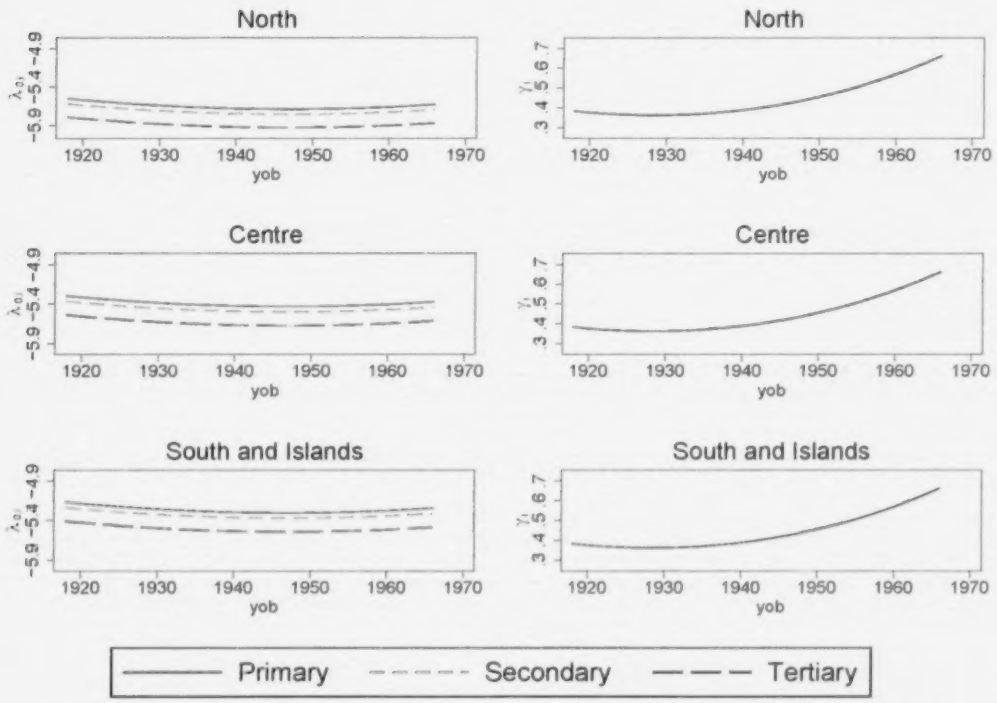

Notes: These figures are generated from Table 3. They show the estimated heterogeneity in the level ($\hat{\lambda}_{0,i}$) and ATM card-holders shift ($\hat{\gamma}_{0,i}$) in money demand.

Figure 6: ATM Card Adoption, Money-Consumption and Wealth-Consumption Ratio

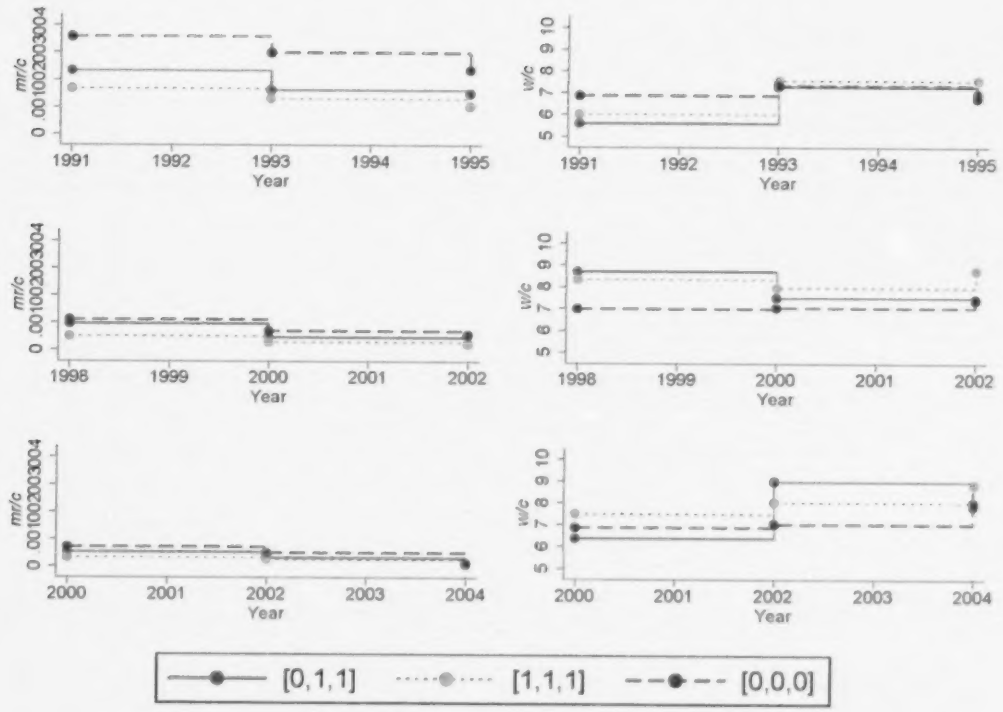

Note: Authors' calculations of the mean money-consumption ratio (mr/c) and wealth-consumption ratio (w/c) over a three-period window for adopters of ATM cards (0,1,1) with adoption during the year: 1993 (top left), 2000 (top right), and 2002 (bottom left). For comparison, the always-adopters of ATM cards (1,1,1) and never-adopters (0,0,0) are provided.

Figure 7: Sensitivity of Adoption-Cost Estimates to the Discount Factor

CRRA Utility

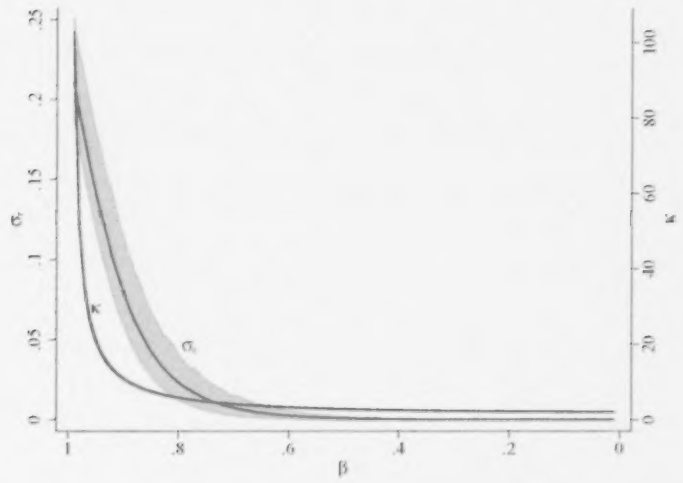

Log Utility

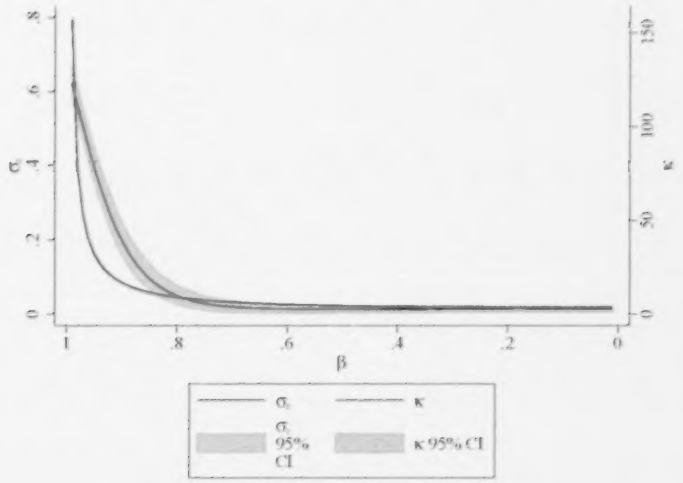

Note: This figure illustrates the median parameter estimates and the 95 percent confidence intervals (based on 1,000 bootstrap samples) of $\tilde{\sigma}$, (blue, left vertical axis) and $\tilde{\kappa}$ (green, right vertical axis) versus the value of the discount factor, β (horizontal axis).